Software Tools for Technology Transfer manuscript No. (will be inserted by the editor)

# A Survey of New Trends in Symbolic Execution for Software Testing and Analysis

Corina S. Păsăreanu<sup>1</sup>, Willem Visser<sup>2</sup>

- Perot Systems Government Services, NASA Ames Research Center, Moffett Field, CA 94035, USA e-mail: Corina.S.Pasareanu@nasa.gov
- <sup>2</sup> SEVEN Networks, Redwood City, CA 94063, USA e-mail: willem@gmail.com

Received: date / Revised version: date

**Abstract.** Symbolic execution is a well-known program analysis technique which represents values of program inputs with symbolic values instead of concrete (initialized) data and executes the program by manipulating program expressions involving the symbolic values. Symbolic execution has been proposed over three decades ago but recently it has found renewed interest in the research community, due in part to the progress in decision procedures, availability of powerful computers and new algorithmic developments. We provide here a survey of some of the new research trends in symbolic execution, with particular emphasis on applications to test generation and program analysis. We first describe an approach that handles complex programming constructs such as input data structures, arrays, as well as multithreading. We follow with a discussion of abstraction techniques that can be used to limit the (possibly infinite) number of symbolic configurations that need to be analyzed for the symbolic execution of looping programs. Furthermore, we describe recent hybrid techniques that combine concrete and symbolic execution to overcome some of the inherent limitations of symbolic execution, such as handling native code or availability of decision procedures for the application domain. Finally, we give a short survey of interesting new applications, such as predictive testing, invariant inference, program repair, analysis of parallel numerical programs and differential symbolic execution.

# 1 Introduction

Modern software systems must be extremely reliable and correct. Automatic methods for ensuring software correctness range from static techniques, such as (software) model checking or static analysis, to dynamic techniques,

such as testing. All these techniques have strengths and weaknesses: model checking is automatic, exhaustive, but may not scale. Static analysis, on the other hand, scales to very large programs but may give too many spurious warnings, while testing alone may miss important errors (since it is inherently incomplete).

We survey here several recent research trends that combine the strengths of these different techniques while overcoming their weakness. In particular, we focus here on approaches to software testing and analysis that are based on symbolic execution. Symbolic execution [12, 36] is a well known program analysis technique that allows execution of programs using symbolic input values, instead of actual data, and represents the values of program variables as symbolic expressions. As a result, the outputs computed by a program are expressed as a function of the symbolic inputs. Its applications range from automated test input generation to proving program partial correctness. Symbolic execution has been proposed over three decades ago but recently it has found renewed interest in the research community, due in part to the progress in decision procedures, availability of powerful computers and new algorithmic developments.

In this paper we begin with a description of our own approach [35,40] to symbolic execution that uses a model checker to explore different symbolic execution paths (Section 2). The approach applies to Java programs and it handles complex input data structures, arrays, as well multi-threading.

Performing symbolic execution on looping programs may result in a large (possibly unbounded) number of symbolic program configurations that need to be analyzed. Therefore symbolic execution might not terminate and in practice, we need to put a limit on the number of such symbolic configurations. An alternative is to use abstraction techniques to try to limit the symbolic space explored during symbolic execution. Our abstrac-

tions are inspired by the ones used in *shape analysis* [38] and are described in Section 3.

We also discuss a popular recent technique (proposed by others) that combines symbolic with concrete execution [28, 46] to overcome some of the inherent limitations of symbolic execution, such as availability of decision procedures and handling calls to native libraries (Section 4). Other related hybrid approaches are discussed in the same section.

We follow with a description of various "classical" applications such as test input and sequence generation, proving program correctness, and static detection of runtime errors. We also describe some novel, "not so classical" applications, that use symbolic execution or its variants for predictive testing, dynamic invariant generation, data structure repair, analysis of parallel numerical programs and differential symbolic execution (Section 5). Section 6 gives a short conclusion.

We give most of our presentation in terms of Java (because this was the context of our own work) but we believe that most of the presentation could also be generalized to other languages.

### 2 Symbolic Execution

#### 2.1 Background

The main idea behind symbolic execution [12, 36] is to use *symbolic values*, instead of actual data, as input values, and to represent the values of program variables as symbolic expressions. As a result, the output values computed by a program are expressed as a function of the input symbolic values.

The *state* of a symbolically executed program includes the (symbolic) values of program variables, a *path condition* (PC) and a program counter. The path condition is a (quantifier-free) boolean formula over the symbolic inputs; it accumulates constraints which the inputs must satisfy in order for an execution to follow the particular associated path. A *symbolic execution tree* characterizes the execution paths followed during the symbolic execution of a program. The tree nodes represent program states and they are connected by program transitions.

Consider the code fragment in Figure 1 (left) [35], which swaps the values of integer variables x and y, when x is greater than y. Figure 1 (right) shows the corresponding symbolic execution tree. Initially, PC is true and x and y have symbolic values X and Y, respectively. At each branch point, PC is updated with assumptions about the inputs, in order to choose between alternative paths. For example, after the execution of the first statement, both then and else alternatives of the if statement are possible, and PC is updated accordingly. If the path condition becomes false, i.e., there is no set of inputs that satisfy it, this means that the symbolic

state is not reachable, and symbolic execution does not continue for that path. For example, statement (6) is unreachable.

# 2.2 Exploring the symbolic execution tree using a model checking tool

Symbolic execution traditionally arose in the context of checking sequential programs with a fixed number of integer variables. Several recent approaches [10, 13, 20] implement dedicated tools to perform various program analyses based on some form of symbolic execution.

In our past work [35] we have defined a generalization of traditional symbolic execution that does not require a dedicated tool but instead enables a standard model checking tool (for the underlying language) to perform symbolic execution. Our approach targets Java programs and it handles complex input data structures and arrays (via "lazy initialization" as explained below) as well as concurrency. The Java PathFinder (JPF) model checking tool [32] is used to explore the symbolic execution tree of the analyzed program. Thus, we take advantage of the model checker's built-in state space exploration capabilities, such as different search strategies (e.g., heuristic search), checking of temporal properties, and partial order and symmetry reductions. A similar tool [19] uses the Bogor model checking framework, instead of JPF, and a "lazier" treatment of initialization for input data structures.

We defined a source-to-source translation that instruments a Java program by adding non-determinism and support for manipulating formulae that represent path conditions in such a way that it enables JPF to perform symbolic execution of the program. The model checker checks the symbolic state space of the program using its usual state space exploration techniques. A symbolic state includes a heap configuration, a path condition on primitive fields, and thread scheduling. Whenever a path condition is updated, it is checked for satisfiability using off-the-shelf decision procedures, such as the Omega library [45] for linear integer constraints. If the path condition is unsatisfiable, the model checker backtracks. Pre-conditions are used to restrict the symbolic search space (to only enable exploration of inputs that satisfy the preconditions).

A specialized type-dependence analysis [2] can be used to minimize the instrumentation effort, by determining which parts of the code depend on the inputs and therefore needs to be instrumented (the rest of the code remaining unchanged). We describe some details of the instrumentation in Section 2.8 (in the context of handling input arrays).

Recently, we have investigated a second approach, that does not require the program instrumentation, but instead implements a non-standard interpreter of Java bytecodes [43].

```
int x, y;
1: if (x > y) {
2:    x = x + y;
3:    y = x - y;
4:    x = x - y;
5:    if (x - y > 0)
6:     assert(false);
}
```




Fig. 1. Code that swaps two integers and the corresponding symbolic execution tree (transitions are labeled with program control points)

# 2.3 Checking safety properties and generating test inputs

Our symbolic execution framework can be used for finding errors to safety properties and for test input generation. Safety properties can be written in the logical formalism recognized by the model checker or they can be specified with code instrumentation, as in [7]. While checking correctness, the model checker reports counterexample(s) that violate a correctness criterion. While generating test inputs, the model checker generates paths that are witnesses to a testing criterion encoded as a safety property (see e.g. [25,31]). For a reported counterexample, the model checker also reports the input heap configuration, the path condition for the primitive input fields thread scheduling, which can be used to reproduce the error.

# ${\it 2.4~Handling~multi-threaded~and~non-deterministic}\\ systems$

As mentioned, our approach allows a standard model checker to perform symbolic execution. We use the model checker also to systematically analyze thread interleavings and other forms of nondeterminism that might be present in the code.

#### 2.5 Loops, recursion, method invocations

We exploit the model checker's search abilities to handle arbitrary program control flow. We do not require the model checker to perform state matching, since state matching is, in general, undecidable when states represent path conditions on unbounded data. Note also that performing (forward) symbolic execution on programs with loops can explore infinite execution trees. Therefore, for systematic state space exploration we put a limit on the search depth of the model checker or we

limit the size of the constraints in the path condition. Note that our symbolic approach can be used for finding counterexamples to safety properties; it can prove correctness for programs that have finite execution trees and have decidable data constraints. For proving properties of programs with unbounded loops, one would need to annotate the program with loop invariants (see discussion in Section 5.3).

#### 2.6 Handling Input Data Structures

We use a *lazy initialization* algorithm for symbolically executing a method that takes as inputs complex data structures with unbounded data. The algorithm starts execution of the method on inputs with *uninitialized* fields and it assign values to these fields "lazily", i.e., when they are first accessed during the method's symbolic execution. This allows symbolic execution of methods without requiring an a priori bound on the number of input objects.

We explain how the algorithm symbolically executes a method with one input object, i.e., the implicit input this. Methods with multiple parameters are treated similarly.

To execute a method m in class C, the algorithm first creates a new object o of class C with uninitialized fields. Next, the algorithm invokes o.m() and the execution proceeds following Java semantics for operations on reference fields and following traditional symbolic execution for operations on primitive fields, with the exception of the special treatment of accesses to uninitialized fields.

When the execution accesses an uninitialized reference field, the algorithm nondeterministically initializes the field to null, to a reference to a new object with uninitialized fields, or to a reference of an object created during a prior field initialization; this systematically treats aliasing. When the execution accesses

- an uninitialized primitive field, the algorithm first initializes the field to a new symbolic value of the appropriate type and then the execution proceeds according to the standard execution semantics.
- When the execution evaluates a branching condition on primitive fields, the algorithm nondeterministically adds the condition or its negation to the corresponding path condition and checks the path condition's satisfiability using a decision procedure. If the path condition becomes infeasible, the current execution terminates (i.e., the algorithm backtracks).

#### 2.7 Example

We illustrate how lazy initialization works using the example from Figure 2 (left), which gives the Java declaration of a class Node that implements singly-linked lists. The fields elem and next represent, respectively, the node's integer value and a reference to the next node. The method swapNode destructively updates its input list (referenced by the implicit parameter this) to sort its first two nodes and returns the resulting list.

We used symbolic execution to check that there are no unhandled runtime exceptions during any execution of swapNode. The result of the check is that the property holds; the analyzed executions are summarized in Figure 2 (right). These executions together represent all possible actual executions of swapNode. For each execution, we show the corresponding input structure, the constraint on the integer values in the input and the output structure. Thus for each row, any actual input list that has the given structure and has integer values that satisfy the given constraint, would result in the given output list. The value "?" for an elem field indicates that the field is not accessed and the "cloud" indicates that the next field is not accessed.

If we comment out the check for null on line (1) in swapNode, our framework reports that for the top most input in Figure 2, the method raises an unhandled NullPointerException. All other input/output pairs stay the same.

The symbolic execution tree in Figure 3 illustrates the (simplified) symbolic execution tree that results from the symbolic execution of swapNode. Each node of the execution tree denotes a *state*, which consists of the state of the heap (including the symbolic values of the elem fields) and the path condition accumulated along the branch (path) in the tree. A transition of the execution tree connects two tree nodes and corresponds to either execution of a statement of swapNode or to a lazy initialization step. Branching in the tree corresponds to a nondeterministic choice that is introduced to handle aliasing or build a path condition.

Symbolic execution starts by first creating a new node object and invoking swapNode on the object. The first access to the uninitialized next field happens at line

(1) and causes it to be initialized. Lazy initialization explores three possibilities: either the field is null or the field points to a new symbolic object or the field points to a previously created object of the same type (with the only option being itself). Intuitively, this means that, at this point in the execution, we make three different assumptions about the configuration of the input list, according to different aliasing possibilities. Another field initialization happens during execution of statement (4), which results in four possibilities, as there are two Node objects at that point in the execution.

When a condition involving primitive fields is symbolically executed, e.g., statement (2), the execution tree has a branch corresponding to each possible outcome of the condition's evaluation. Evaluation of a condition involving reference fields does not cause branching unless uninitialized fields are accessed.

If swapNode has the precondition that its input should be acyclic, then symbolic execution does not explore the transitions marked with an "X".

In order to keep track of the input data structures for programs with *destructive updating*, we build mappings between objects with un-initialized fields and objects that are created when those fields are initialized (these maps are used to re-construct the input structures, e.g. for test input generation).

#### 2.8 Handling Input Arrays

Symbolic execution for programs that have as inputs arrays of unspecified size can also use lazy initialization [40].

Consider the code shown in Figure 4 (left). This method takes as a parameter an array of integers a and it sets all the elements of a to zero. This method has a precondition that its input is not null. The assert clause declares a partial correctness property that states that after the execution of the loop, the value of the first element in a is zero (we will describe in Section 5.3 how we can use symbolic execution and loop invariants to prove this property).

In order to symbolically execute the code we first instrument it to enable JPF to perform symbolic execution. The instrumented code and part of the library classes that we provide are illustrated in Figure 4 (right) and Figure 5, respectively. The interested reader is referred to [35, 40] for a detailed description of code instrumentation, here we just highlight some key features.

The main idea is to replace concrete types with corresponding "symbolic types" (i.e. library classes that we provide) and concrete operations with method calls that implement "equivalent" operations on symbolic types. Classes Expression and IntArrayStructure support manipulation of symbolic integers and symbolic integer arrays, respectively. The static field Expression.pc stores the (numeric) path condition. Method \_update\_LT makes a nondeterministic choice

```
class Node {
                                                                   Input list
                                                                                       Constraint =>
                                                                                                        Returned list
  int elem;
                                                                      ?
                                                                                                           ?
                                                                                          none
  Node next;
                                                                     E0 next
                                                                                          none
  Node swapNode() {
1: if(next!=null)
                                                                  E0 next E1 next
                                                                                       (E0 <= E1)
                                                                                                       E0 next E1 next
    if(elem-next.elem>0){
3:
       Node t = next;
                                                                                        (E0 > E1)
4:
       next = t.next;
5:
       t.next = this;
                                                                                        (E0 > E1)
6:
       return t;
                                                                                        (E0 > E1)
7: return this;
                                                                                        (E0 > E1)
}
```

Fig. 2. Code to sort the first two nodes of a list (left) and an analysis of this code using our symbolic execution based approach (right)




Fig. 3. Symbolic execution tree (excerpts)

```
void example() {
// @ precondition: a != null;
                                                IntArrayStructure a = new IntArrayStructure();
void example(int[] a) {
                                                Expression i = new IntegerConstant(0);
    int i = 0;
2:
    while (i < a.length) {
                                                while(Expression.pc._update_LT(i,a.length)) {
3:
      a[i] = 0;
                                                  a._set(i,new IntegerConstant(0));
4:
                                                   i = i._plus(new IntegerConstant(1));
    }
   assert a[0] == 0;
                                                 assert Expression.pc._update_EQ(
                                                        a._get(new IntegerConstant(0)),0);
```

Fig. 4. Array example (left) and corresponding instrumented code (right)

```
class Expression { ...
                                                         class IntArrayStructure {
 static PathCondition pc;
                                                          Vector _v;
                                                          Expression length;
Expression _plus(Expression e){
   ...} }
                                                          ArrayCell _new_ArrayCell(Expression idx) {
                                                           for(int i=0;i<_v.size();i++) {</pre>
class PathCondition { ...
                                                            ArrayCell cell=(ArrayCell)_v.elementAt(i);
 Constraints c;
                                                            if(Expression.pc._update_EQ(cell.idx,idx))
boolean _update_LT(Expression 1,
                                                             return cell;
                     Expression r){
                                                           ArrayCell t=new ArrayCell(length,idx,name);
  boolean result;
                                                           _v.add(t);
  result=Verify.choose_boolean();
                                                           return t;
  if (result)
    c.add_constraint_LT(e1,e2);
                                                          public Expression _get(Expression idx) {
  else
                                                           assert(Expression.pc._update_GE(idx, 0)&&
    c.add_constraint_GE(e1,e2);
                                                                Expression.pc._update_LT(idx,length));
                                                           ArrayCell cell = _new_ArrayCell(idx);
  Verify.ignoreIf(!c.is_sat());
                                                           return cell.elem;
  return result;
  } }
```

Fig. 5. Library classes

(i.e., a call to choose\_boolean) to add to the path condition the constraint or the negation of the constraint its invocation expresses and returns the corresponding boolean. Method is\_sat uses a decision procedure to check if the path condition is infeasible (in which case, JPF will backtrack). Method \_plus constructs a new Expression that represents the sum of its input parameters. IntegerConstant is a subclass of Expression and wraps concrete integer values.

To store the input array elements that are created as a result of a lazy initialization, we use a variable of class Vector, for each input array. The \_get and \_set methods use the elements in this vector to systematically initialize input array elements. When the execution accesses a symbolic array cell, the algorithm nondeterministically initializes it to a new cell or to a cell that was created during a prior cell initialization. The assertion checks in the \_get/\_set methods establish that there are no array out of bounds errors.

#### 2.9 Other Challenges to Symbolic Execution

Other typical challenges to symbolic execution include handling common library classes and/or native code (i.e. code that can not be analyzed directly by symbolic execution). Such code needs to be modeled explicitly to be considered by the symbolic execution (see e.g. [44]).

A promising approach that targets Java *string* library classes is presented in [47]. In that work, the implementation details of strings are abstracted away the using finite state automata, resulting in scaling of symbolic execution to complex string manipulating applications.

Section 4 describes an orthogonal technique that combines concrete and symbolic execution to address this problem.

#### 2.10 Integrating Multiple Decision Procedures

Perhaps the main challenge to symbolic execution is the availability of the decision procedures for the application domain and the number of constraints that can be handled by the decision procedure/constraint solver. This is why symbolic execution is most effective at unit or subsystem level; i.e. for analyzing a procedure or a set of procedures.

To (partially) alleviate this problem, we equipped our symbolic execution framework with a generic interface to multiple decision procedures (e.g. CVC3, Yices, STP, etc., [4]). More recently, we have also integrated two constraint solvers for real constraints (Choco and IASolver) [43]).

The user can chose between multiple decision procedures that interact in different modes with the symbolic execution framework. Furthermore, there are different optimizations possible for this interaction, e.g. if the decision procedure supports incremental constraint solving, the path condition is not sent at once to the constraint solver for solving, but rather just the new constraint that should be added/removed before checking satisfiability.

#### 3 Abstraction

As mentioned, performing symbolic execution on looping programs may result in an infinite execution tree. Therefore we perform search with limited depth, or put a limit on the number of constraints in the path condition. An alternative approach [3,52] considers state matching techniques to limit the state space search. The approach involves checking when a symbolic state  $(s_i)$  is subsumed by another symbolic state  $(s_j)$ , i.e., the set of

concrete states represented by  $s_i$  is included in the set of concrete states represented by  $s_i$ .

Subsumption is used to determine when a symbolic state is revisited, in which case the model checker backtracks, thus pruning the state space search. Even with subsumption, the number of symbolic states may still be unbounded. We therefore defined abstraction mappings to be used during state matching. More precisely, for each explored state, the model checker computes and stores an abstract version of the state, as specified by the abstraction mappings. Subsumption checking then determines if an abstract state is being revisited. This effectively explores an under-approximation of the (feasible) paths through the program. Therefore the technique is still useful for finding safety errors or for test input generation (see Section 5.2 for a discussion of applications of abstract subsumption in the context of test sequence generation).

#### 3.1 Example

In our approach [3] we defined abstract subsumption checking for singly linked lists and arrays (by reducing their representation to lists). The abstraction that we have implemented are inspired by [38,57] and are based on the idea of summarizing all the nodes in a maximally uninterrupted list segment with a summary node. The main difference between [38,57] and our abstractions is that we also summarize the numeric data stored in the summarized nodes and we give special treatment to uninitialized nodes. The numeric data stored in the abstracted list is summarized by setting the valuation for the summary node to be a disjunction of the valuations of the summarized nodes. Intuitively, the numeric data stored in a summary node can be equal to that of any of the summarized nodes.

We illustrate abstract subsumption for singly-linked lists using the example in Figure 6. For more details, please see the related paper [3].

Figure 6 depicts two symbolic states,  $s_8$  and  $s_{12}$  that resulted during the analysis of a list manipulating program [3]. These states can not be matched, since their "heap shape" is different. However, let us consider the abstract heap shape and the corresponding valuations for state  $s_{12}$ . The abstracted state is subsumed by state  $s_8$  since the corresponding heap shapes match (as illustrated by the common node labels  $l_1$ ,  $l_2$ ,  $l_3$ ). Furthermore, there is a valid logical implication between the normalized numeric constraints of the two states.

#### 4 Combining Concrete and Symbolic Execution

Several recent tools implement a new hybrid analysis approach, that performs a concrete execution along symbolic execution for dynamic test generation, e.g.

```
1: void foo(int x,int y){
2:    int z = x*x*x; /* could be z = h(x) */
3:    if (y == z) {
4:       assert(false); /* error */
5:    }
6: }
```

Fig. 7. Code for illustrating concolic execution


DART [28], CUTE [37,46], EXE [11], PEX [42]. This popular approach has been applied to finding errors in many challenging areas such as Web and data-base applications [6,21,54].

The idea is to perform a concrete execution on random inputs and at the same time to collect the path constraints along the executed path; this is also called "concolic execution". These path constraints are then used to compute new inputs that drive the program along alternative paths. More specifically, one can negate one constraint at a branch point to guide the test generation process towards executing the other branch. An off-the-shelf constraint solver is called to solve the path constraints and to obtain the test inputs. The program is executed on these new inputs, constraints are collected along the new program path and the process is repeated until all the execution paths are covered (therefore it may never terminate) or until the desired test coverage is achieved. The approach works by code instrumentation and does not use model checking (therefore can not analyze multithreading systematically). However, the main advantage of this hybrid approach is that the concrete execution can be used "to help" the symbolic execution in certain situations, e.g. when there are no available decision procedures or in the presence of native calls.

#### 4.1 Example

As an example, consider the code in Figure 7 [26]. Assume we have decision procedures/constraint solvers that can reason about linear constraints only. Initially the inputs that were randomly generated are x = 3 and y = 7. The concrete value of z is 27, but the symbolic value is z = X\*X\*X, and the path condition (corresponding to the else branch) is Y != X\*X\*X; therefore the decision procedures cannot handle it. However, instead of taking the symbolic value z = X\*X\*X in the path condition, one can take the concrete value (i.e. z = 27). The path condition then becomes Y != 27 and the execution continues until the end of the procedure. In order to obtain inputs that guide the execution towards the than branch, one needs to solve Y == 27 which can be done easily with the available constraint solver. The program is then re-executed with the new inputs: x = 3 and y = 27 and the error at line 4 is discovered.

Assume now that instead of int z = x\*x\*x;, statement 2 is int z = h(x);, where h is some library function (alternatively assume its code is simply un-available



**Fig. 6.** Abstract subsumption between  $s_8$  and  $s_{12}$ 

to symbolic execution, e.g., could not be instrumented). Then the same reasoning as above can be applied (therefore eliminating the need for explicit modeling of h). Of course, there may be some situations when such an approach would not be recommended, due to certain side-effects of method h (e.g., writing data to a file that is later read and affects the execution). In that case, some modeling would still be required.

### 4.2 Compositional Symbolic Execution

The main obstacle for scaling hybrid concrete-symbolic execution to reasoning about complex programs is the large (possible infinite) number of paths that need to be explored. Recent work [1, 27] proposes compositional reasoning as a means of scaling up symbolic execution. The work has been done in the context of the hybrid concrete-symbolic execution described above, but we believe that it can also be extended to "classical" symbolic execution (introduced in Section 2).

The idea is to use "summaries" of individual functions (similar to inter-procedural static analysis); these summaries are computed "top down", on a demand driven basis. If f() calls g(), one can analyze/test g() separately, summarize the results, and use g()'s summaries when analyzing/testing f(); thus, each method is analyzed separately and the over-all number of analyzed paths is smaller than in the case the two procedures are analyzed as a whole.

#### 4.3 Other Combined Analyses

In concolic execution the idea is to perform a concrete execution together with a symbolic analysis that is used to produce inputs to cover "new" behavior with the aim to uncover errors. One can also take the opposite approach by first doing a symbolic (usually in-precise) analysis to find a possible error and then perform a concrete execution (i.e. run the program) to determine if it is real or not. The reason for this second step is that the symbolic execution can be unsound (it might follow paths in the code that are not possible in reality); this may happen if the analysis is only intra-procedural (don't follow procedure calls) and just returns new unconstrained symbolic values for the returned values of the procedures

that are not analyzed. The Check&Crash system [16] uses ESC/Java [23] to do the symbolic analysis and then JCrasher to execute the test to see if it is a real test.

In [51] a custom symbolic execution is used that allows inter-procedural analysis in which the degree of procedure nesting can be varied (see Section 5.5 for more details).

Other related hybrid techniques include the use of concrete execution to effectively "set-up" the environment for symbolic execution [44] and a combination of test case generation based on symbolic execution and run-time monitoring [5]; both these techniques have been applied in the context of NASA software systems. Furthermore, other related approaches [29, 58] seek to combine abstraction techniques (with automatic abstraction refinement) and theorem proving for program analysis and testing.

#### 5 Applications/Analyses

Symbolic execution has many applications, most notably in testing and proving program correctness. We discuss them below, together with some exciting new applications.

#### 5.1 Test Input Generation

Obtaining high coverage is always the goal of testing, but the reality is that structural coverage is the only meaningful measure of test adequacy and as such obtaining high structural coverage is often the goal of test case generation techniques. Symbolic execution lends itself particularly well to this task, since the path condition to reach a branch or statement in the code (the two most often used forms of structural coverage in industry; statement and branch coverage) when solved, gives exactly the inputs to reach the statement or branch (i.e. the test inputs for the test case). We refer to this as test generation for white-box testing.

In addition one can also do test generation in a black-box fashion by essentially using the same general technique, but now instead of symbolically executing the program under test, one executes a Java predicate characterizing all valid input structures for the code (often called the representation invariant, or, rep0k() method [9,53]). The objective here is to generate "symbolic" structures that satisfy the representation invariant that can be concretized (by solving the path condition to reach a valid structure) to a valid input for the program under test. This general approach, although not using symbolic execution, was popularized by the Korat tool [9]. See [53] for a detailed description of using symbolic execution to generate test inputs in this fashion.

### 5.2 Test Sequence Generation

Both the white- and black-box techniques described above suffer from the issue that one can generate inputs that are actually not possible during normal execution of the program. In the white-box case this can happen since it is typical to analyze each API (application programming interface) call for a system in isolation and it may happen that in reality the calling context of a method may provide some implicit pre-conditions. Similarly in the black-box case it may simply be that although a certain input is legal it can actually never be provided as an input (i.e. it can not be constructed using the public methods/fields allowed by the respective java class).

To alleviate these concerns one can generate sequences of inputs, rather than single input methods [52, 56]. As a simple example, consider a class BinTree that provides a Java implementation of binary search trees.

```
public class BinTree {
  private Node root;
  ...
  public void add (int x) { ...}
  public boolean remove (int x) { ... }
}
```

A test sequence for this class is as follows:

```
BinTree t = new BinTree();
t.add(1); t.add(2); t.remove(1);
```

It contains a sequence of method calls in the class interface (e.g. add and remove), with some method arguments, that builds relevant object states and exercise the code in some desired fashion (e.g. to achieve statement or predicate coverage [52]).

Generating test sequences can be done by enumerating all the possible test sequences (up to a given size) and executing them symbolically (to account for the method arguments). The main problem now however becomes that analyzing all combinations of possible interface calls quickly produces a state explosion. The solution is to provide a mechanism for state-matching between API calls in this symbolic case.

Although this problem is undecidable in general, if one only considers container classes storing integer data (a very common case) it is tractable. One can also match states using an abstraction of the state (as explained in Section 3), i.e. match abstract versions of states where the concretized states will not match. The trade-offs are obvious, match too liberally (i.e. using abstraction) and the coverage will not be obtained, and match too finely (i.e. check full subsumption on symbolic states), and run the risk of never terminating the search.

Using the shape of the container as the abstraction function was found to be particularly powerful [52]: for example, we could show that the shortest sequence of API calls on a Fibonacci Heap implementation to obtain statement coverage was 12. This is an interesting result in itself, since the code is only a few hundred lines long and to obtain the simplest form of coverage requires 12 calls.

For a detailed study of the various techniques for generating test sequences for container classes see [52] (all examples are made available though the JPF Source-Forge website). We analyzed java implementations for Binary Tree, Fibonacci Heap, Binomial Heap, Tree Map). We compared explicit state model checking, symbolic and concrete execution (with and without abstract matching) and random testing. We found that symbolic execution worked better than explicit model checking and that, not surprisingly, shape abstraction provides an accurate representation of containers. We found that random testing worked pretty well but it requires longer sequences to achieve good coverage.

#### 5.3 Proving Program Properties

If there is an upper bound on the number of times each loop in the program may be executed, symbolic execution can be used for proving correctness, since the corresponding symbolic execution tree is finite.

However, for most programs, no fixed bound on the number of times each loop is executed exists and the corresponding execution trees are infinite. In order to prove the correctness of such programs, one needs traversing the symbolic execution tree inductively rather than explicitly [30], using annotations in the form of loop invariants. Such annotations are provided by the user or may be discovered automatically, see e.g. [14,15,24,39,40,49,55]. Recent tools that implement such reasoning include ESC/Java [23] (it does not use traditional symbolic execution, but similar symbolic reasoning) and Bogor/Kiasan [19] for reasoning about properties of Java programs. Furthermore, Smallfoot [8] uses symbolic execution and separation logic for proving Hoare-style triples on heap-manipulating programs.

For simplicity of presentation, we illustrate the technique on a single-loop program such as the one in Figure 8 (left); multiple loops can be treated similarly, see e.g. [55]. The program consists of some (loop-free) initialization code, a loop with condition C and (loop-free) body B, and post condition P.

To verify that P holds, it suffices to find a loop invariant I, i.e. a formula that is true when entering the

```
init;
assert I; /* base case */
make symbolic variables read in B;
assume I;
if (C) {
    B;
}
assert P;
}
else
assert P;
```

Fig. 8. Single loop program (left) and instrumented program for proof (right)

loop, re-entering the loop during its iteration and exiting the loop [30]. Moreover, I must be strong enough to produce verifiable results (hence a loop invariant true is, in general, not sufficient). In a symbolic execution framework, this amounts to checking the three assertions in the modified program in Figure 8 (right). Here, we replaced the while statement with an if statement; this is equivalent to placing a "cut" in the loop [30]. At this cut point, we consider all the variables that are modified in the loop body initialized to new symbolic values, and the path condition initialized to true. Note that a symbolic execution from this point on is representative of an arbitrary number of loop unrollings; the "input variables" at the cut point are the variables that are modified by the loop body and their new symbolic values represent all cases. Since the program loop has been cut, this symbolic execution will terminate and have a finite symbolic execution tree.

We then use symbolic execution to check three assertions :

- the assertion at line (4) is the *base case* of the inductive argument and checks that *I* holds when entering the loop
- the assertion at line (7) is the induction step and checks that, assuming I holds at the beginning of the loop, I also holds after the execution of the loop body (i.e. I is inductive)
- the assertion at line (9) checks that I is strong enough for the property to hold (i.e.  $I \land \neg C \to P$ )

If there are no assertion violations in the loop-free program of Figure 8 (right), then the program of Figure 8 (left) does not violate the property P.

#### 5.4 Example

As an example, consider again the code presented in Figure 4. Using the loop invariant  $i \geq 0$ , symbolic execution can be used to automatically check that there are no array bounds violations. This is a simple invariant that can be stated without much effort. In order to prove that there are no assertion violations, a more complex loop invariant is needed, namely  $\neg(a[0] \neq 0 \land i > 0)$ . In [40] we present a technique that generates such invariants automatically, by iterative approximation. The technique

handles different types of constraints (e.g. boolean or numeric, constraints on dynamically allocated data and arrays) and it allows for checking universally quantified formulas. Such formulas are necessary for expressing properties of programs that manipulate unbounded data (such as the input array in Figure 4)

#### 5.5 Static Detection of Run-time Errors

Using symbolic execution to find potential runtimeerrors is a well-known technique. The most famous example of this is the success of Intrinsa's PREfix tool [10] that ultimately led to a buy-out by Microsoft. More recent examples include the work of Engler et al. in [11] for detecting runtime errors in C code and Tomb et al. in [50] that detects errors in Java code.

The idea behind all these tools is to symbolically execute a program until a state is reached where a runtime violation is "possible", for example a null-pointer dereference, division by zero, etc., and a potential error is reported. Unfortunately, due to mostly scalability issues, one can often not execute programs from their inputs, thus it is common to only analyze public or API methods and often times only intra-procedurally. This means the analysis can report errors that are not possible, so-called spurious errors.

One approach to reduce the possible false positives is to use the "variably inter-procedural" analysis described in [50]. As the name suggests the idea here is to allow one to vary the level of the inter-procedural analysis to follow calls n levels deep. Furthermore the approach proposes to solve the input constraints that are associated with a possible error and to form a test case; the analysis reports the error only if the test case actually produces the expected error (similar to Check-n-Crash [16]).

#### 5.6 Examples

As an illustration of some of the advantages of variably inter-procedural analysis, consider the program in Figure 9 and the problem of detecting null pointer dereferences. Lets first assume we use an intraprocedural analysis where we don't follow the calls to the Integer.toHexString method (as is done in [16]);

```
1: class Example {
     public String hexAbs(int x) {
2:
       String result = null;
3:
4:
       if(x > 0)
5:
         result = Integer.toHexString(x);
6:
       else if (x < 0)
7:
         result = Integer.toHexString(-x);
       return result.toUpperCase();
8:
9:
     }
10:}
```

Fig. 9. A simple Java program that illustrates some benefits of symbolic execution.

```
1: int target = ...;
2: int delta = ...;
3: foo(int i) {
4:
     if (similar(i,target)) {
       y = 10/i; // interesting code
5:
6:
7: }
8: ...
9: boolean similar (int i, int target) {
10:
    if (((target - delta) <= i) &&
          (target + delta) >= i)
11:
12:
       return true;
13:
    return false;
14:}
```

Fig. 10. An example where intra-procedural analysis is sufficient.

a possible null pointer dereference will be flagged at line 8, with no constraints on the value of x.

Using variably inter-procedural symbolic execution, we can do better. If we set the analysis to evaluate all method calls up to a depth of 1, it can follow the calls to Integer.toHexString, and determine that they never return null values. Then, because it is a path-sensitive analysis, it can determine that a null pointer dereference can only happen (and must happen) if x=0. Thus, the analysis has ruled out the false positives (the assignments on lines 5 and 7), and has given more information about the true error (the missing case for x=0). Given the constraint on x, it is then straightforward to construct a test case that will trigger the bug.

Varying the level of inter-procedural analysis can have some interesting consequences, for example in [50] it was found that going from an intra-procedural to an inter-procedural analysis might not find more errors but will reduce the number of possible errors the symbolic analysis discovers (and thus will lead to test cases to run to see if it is a real error). The code in Figure 10 illustrates the intuition for this behavior. Note that depending on the value of target and delta there could be a division by zero in this code. Let's assume we pick target = 100 and delta = 10, in which case there is no division by zero. The result of an intra-procedural analysis is one warning, but no error (since the warning corresponds to the case when i = 0 and that would make the division

```
1: foo(int m) {
2:    answer(m);
3:    m = m/(1-m);
4: }
5: ...
6: int answer(int v) {
7:    return v == 42 ? 1: 0;
8: }
```

Fig. 11. An example where inter-procedural analysis is required.

unreachable). The reason for this behaviour is that the call to *similar* is ignored and a fresh symbolic variable is created to hold the result of the call.

However, an inter-procedural analysis results in no warnings (and no errors) since the constraints on similar combined with the fact that i is 0 become infeasible.

The interesting case here is if we pick the values to expose the problem (e.g. change target to 1). Now both an intra- and an inter-procedural analysis expose the error. Note that an intra-procedural analysis also finds the problem simply because the statement is reachable (by picking target and delta to expose the problem); thus adding the constraint that i should be 0 to have a possible division by zero is enough to actually find the error.

One can also create an example to show the opposite effect where obtaining additional constraints actually exposes errors that would otherwise not have been found — this happens when analyzing the code in Figure 11. Here an intra-procedural analysis has no additional constraints on the input value m and thus the chances of the test generation to randomly pick 42 is almost zero. However during an inter-procedural analysis the constraint that m should be 42 is recorded and that would make picking m trivial to expose the division by zero error.

In general a statement that is potentially buggy can be reached in many more ways that don't expose the error than in ways that will expose the error — if this is not true then the error will be found and fixed quickly anyways. Therefore the additional constraints one obtains by doing an inter-procedural analysis will mostly reduce the number of infeasible paths (of an intra-procedural analysis) that reach a potentially buggy statement but it will not necessarily increase the likelihood of generating a test to reach the error.

An enhancement to the general approach of symbolic execution for finding runtime errors is suggested in [22] where it is pointed out that the analysis can be optimized by taking the unconstrained inputs to a program and then constraining them by the negation of the path conditions corresponding to paths that lead to errors. For example, consider the following code:

```
public void foo(Object o) {
  o.x = 5;
  ...
}
```

Assume o is unconstrained; a possible null-pointer exception will be flagged on the dereference in the first line. However since o is unconstrained one ignores this error and rather remove the unconstrained "tag" from o and replace it with the constraint that o is from now on non-null. This technique eliminates false positives and in addition constrains possible executions which allows better scaling. One can ask but what if o was really null? To account for that one can simply rank these potential errors as lower priority to consider than ones that are obtained from using the suggested technique.

## 5.7 Other Applications

Symbolic execution has many applications and it is impossible to enumerate them all. We can only list here a few new "not so standard" applications of symbolic execution (and related hybrid approaches):

- Predictive Testing [33] attempts to predict errors from correct traces. The idea is to perform a "concolic execution" along concrete traces generated by running an existing test suite and to check for assertion violations and other types of errors along these executions: the assertions that hold along a concrete execution do not necessarily hold along the corresponding symbolic execution (since the latter characterizes multiple concrete executions).
- Invariant Inference [17] generates "likely" program invariants in the form of method pre- and post- conditions and class invariants that hold for a given set of tests; the technique is similar in spirit to Daikon [18] but uses the constraints collected during a symbolic execution to come up with the invariants, instead of the invariant patterns used by Daikon.
- Program and Data Structure Repair can be done using symbolic execution; e.g., given an assertion that represents desired structural integrity constraints and a structure that violates them, the algorithm from [34] can "mutate" the given structure to satisfy the constraints.
- Parallel Numerical Program Analysis [48] involves combining model checking and symbolic execution to establish the equivalence of a sequential and a parallel program. The sequential program acts as the "specification" for the parallel one. The symbolic execution is particularly tailored to handling floating point arithmetic.
- Differential Symbolic Execution [41] computes the "logical" differences between two versions of a program; such differences can be used to automate software evolution tasks such as regression test maintenance, reducing re-certification activities or checking behavioral equivalence of two programs after software re-factoring.

#### 6 Conclusions

In this paper, we surveyed new techniques based on symbolic execution and we discussed some of their "traditional" applications, such as test generation and program analysis, as well as some new, interesting applications. The work related to the subject here is vast and it is simply impossible to cover it all in one article. However, we hope that this survey (albeit very limited) will serve as a starting point for more new, exciting applications in this area. For instance, an avenue for immediate future research would be to "parallelize" all/any of the analyses presented in this article.

#### References

- S. Anand, P. Godefroid, and N. Tillmann. Demanddriven compositional symbolic execution. In Proc. of TACAS, 2008.
- S. Anand, A. Orso, and M. J. Harrold. Type-dependence analysis and program transformation for symbolic execution. In Proc. of the 13th International TACAS Conference, 2007.
- 3. S. Anand, C. S. Păsăreanu, and W. Visser. Symbolic execution with abstract subsumption checking. In *Proc.* of the 13th International SPIN Workshop, 2006.
- S. Anand, C. S. Păsăreanu, and W. Visser. JPF-SE: A symbolic execution extension to Java PathFinder. In Proc. of the 13th International TACAS Conference, 2007.
- C. Artho, H. Barringer, A. Goldberg, K. Havelund, S. Khurshid, M. R. Lowry, C. S. Pasareanu, G. Rosu, K. Sen, W. Visser, and R. Washington. Combining test case generation and runtime verification. *Theor. Comput. Sci.*, 336(2-3):209-234, 2005.
- S. Artzi, A. Kiezun, J. Dolby, F. Tip, D. Dig, A. Paradkar, and M. D. Ernst. Finding bugs in dynamic web applications. In *Proc. of ISSTA*, 2008.
- T. Ball, R. Majumdar, T. Millstein, and S. Rajamani. Automatic predicate abstraction of C programs. In Proceedings of the 2001 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI 2001), volume 36.5 of ACM SIGPLAN Notices, pages 203-213. ACM Press, June 2001.
- 8. J. Berdine, C. Calcagno, and P. O'Hearn. Symbolic execution with separation logic. In *Proc. of Third Asian Symposium*, 2005.
- C. Boyapati, S. Khurshid, and D. Marinov. Korat: Automated testing based on Java predicates. In *International Symposium on Software Testing and Analysis (IS-STA'02)*, pages 123–133, July 2002.
- 10. W. R. Bush, J. D. Pincus, and D. J. Sielaff. A static analyzer for finding dynamic programming errors. *Software: Practice and Experience*, 30(7):775–802, 2000.
- C. Cadar, V. Ganesh, P. M. Pawlowski, D. L. Dill, and D. R. Engler. EXE: automatically generating inputs of death. In Proc. ACM Conference on Computer and Communications Security, 2006.
- 12. L. A. Clarke. A system to generate test data and symbolically execute programs. *IEEE Trans. Software Eng*, 2(3):215–222, 1976.

- 13. A. Coen-Porisini, G. Denaro, C. Ghezzi, and M. Pezze. Using symbolic execution for verifying safety-critical systems. In Proceedings of the 8th European software engineering conference held jointly with 9th ACM SIGSOFT international symposium on Foundations of software engineering, pages 142-151. ACM Press, 2001.
- M. Colon, S. Sankaranarayanan, and H. Sipma. Linear invariant generation using non-linear constraint solving. In Proc. CAV, 2003.
- P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In Proc. 5th POPL, 1978.
- C. Csallner and Y. Smaragdakis. Check 'n' crash: Combining static checking and testing. In *Proc. of ICSE*, 2005.
- C. Csallner, N. Tillmann, and Y. Smaragdakis. DySy: Dynamic symbolic execution for invariant inference. In Proc. of ICSE, 2008.
- 18. The Daikon invariant detector. http://groups.csail.mit.edu/pag/daikon//.
- X. Deng, J. Lee, and Robby. Bogor/kiasan: A k-bounded symbolic execution for checking strong heap properties of open systems. In *Proc. of ASE*, 2006.
- D. L. Detlefs, K. R. M. Leino, G. Nelson, and J. B. Saxe. Extended static checking. Research Report 159, Compaq Systems Research Center, 1998.
- M. Emmi, R. Majumdar, and K. Sen. Dynamic test input generation for database applications. In *Proc. of ISSTA*, 2007
- D. Engler and D. Dunbar. Under-constrained execution: making automatic code destruction easy and scalable. In International Symposium on Software Testing and Analysis (ISSTA), July 2007.
- C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. In Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation, pages 234-245, 2002.
- C. Flanagan and S. Qadeer. Predicate abstraction for software verification. In Proc. POPL, 2002.
- 25. A. Gargantini and C. Heitmeyer. Using model checking to generate tests from requirements specifications. In Proc. of the 7th European Software Engineering Conference, Held Jointly with the 7th ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE), 1999.
- P. Godefroid. Software model checking via static and dynamic program analysis. In MOVEP.
- 27. P. Godefroid. Compositional dynamic test generation. In  $Proc.\ of\ POPL,\ 2007.$
- P. Godefroid, N. Klarlund, and K. Sen. DART: directed automated random testing. In Proc. PLDI, 2005.
- 29. B. S. Gulavani, T. A. Henzinger, Y. Kannan, A. V. Nori, and S. K. Rajamani. SYNERGY: a new algorithm for property checking. In *Proc. of SIGSOFT FSE*, 2006.
- 30. S. L. Hantler and J. C. King. An introduction to proving the correctness of programs. *ACM Comput. Surv.*, 8(3):331–353, 1976.
- H. Hong, I. Lee, O. Sokolsky, and H. Ural. A temporal logic based theory of test coverage and generation. In Proceedings of the 8th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS), Apr. 2002.

- 32. Java PathFinder. http://javapathfinder.sourceforge.net.
- 33. P. Joshi, K. Sen, and M. Shlimovich. Predictive testing: Amplifying the effectiveness of software testing (short paper). In *Proc. of ESC/FSE*, 2007.
- S. Khurshid, I. Garcia, and Y. Suen. Repairing structurally complex data. In Proc. of SPIN, 2005.
- S. Khurshid, C. S. Păsăreanu, and W. Visser. Generalized symbolic execution for model checking and testing. In Proc. of the 9th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), 2003.
- 36. J. C. King. Symbolic execution and program testing. Commun. ACM, 19(7):385-394, 1976.
- R. Majumdar and K. Sen. Hybrid concolic testing. In Proc. of ICSE, 2007.
- R. Manevich, E. Yahav, G. Ramalingam, and M. S. iv. Predicate abstraction and canonical abstraction for singly-lin ked lists. In *Proc. VMCAI*, Paris, 2005.
- 39. Z. Manna and A. Pnueli. The Temporal Logic of Reactive and Concurrent Systems: Specification. 1992.
- 40. C. Pasareanu and W. Visser. Verification of java programs using symbolic execution and invariant generation. In Proceedings of the 11th International SPIN Workshop on Model Checking of Software, volume 2989 of Lecture Notes in Computer Science. Springer-Verlag, 2004.
- S. Person, M. B. Dwyer, S. Elbaum, and C. S. Păsăreanu. Differential symbolic execution. In ACM SIGSOFT 2008/FSE 16, Nov. 2008.
- 42. PEX: Automated Exploratory Testing for .NET. http://research.microsoft.com/Pex/.
- C. S. Păsăreanu, P. Mehlitz, D. Bushnell, K. Gundy-Burlet, M. Lowry, S. Person, and M. Pape. Combining unit-level symbolic execution and system-level concrete execution for testing nasa software. In *Proc. of ISSTA*, 2008.
- 44. C. S. Păsăreanu, P. C. Mehlitz, D. H. Bushnell, K. Gundy-Burlet, M. Lowry, S. Person, and M. Pape. Combining unit-level symbolic execution and systemlevel concrete execution for testing nasa software. In International Symposium on Software Testing and Analysis (ISSTA), July 2008.
- W. Pugh. The Omega test: A fast and practical integer programming algorithm for dependence analysis. Communications of the ACM, 31(8), Aug. 1992.
- K. Sen, D. Marinov, and G. Agha. CUTE: a concolic unit testing engine for C. In Proc. ESEC/SIGSOFT FSE, 2005.
- D. Shannon, S. Hajra, A. Lee, D. Zhan, and S. Khurshid. Abstracting symbolic execution with string analysis. In Proc. of Testing: Academic and Industrial Conference— Practice and Research Techniques (TAIC-PART 2007), 2007.
- S. F. Siegel, A. Mironova, G. S. Avrunin, and L. A. Clarke. Using model checking with symbolic execution to verify parallel numerical programs. In *Proc. of ISSTA*, 2006.
- A. Tiwari, H. Rues, H. Saidi, and N. Shankar. A technique for invariant generation. In Proc. TACAS, 2001.
- A. Tomb, G. Brat, and W. Visser. Variably interprocedural program analysis for runtime error detection. In International Symposium on Software Testing and Analysis (ISSTA), July 2007.

- 51. A. Tomb, G. P. Brat, and W. Visser. Variably interprocedural program analysis for runtime error detection. In *Proc. of ISSTA*, 2007.
- W. Visser, C. S. Pasareanu, and R. Pelanek. Test input generation for java containers using state matching. In *Proc. ISSTA*, 2006.
- 53. W. Visser, C. S. Păsăreanu, and S. Khurshid. Test input generation in Java Pathfinder. In *Proc. of the International Symposium on Software Testing and Analysis (ISSTA)*, 2004.
- 54. G. Wassermann, D. Yu, A. Chander, D. Dhurjati, H. Inamura, and Z. Su. Dynamic test input generation for web applications. In *Proc. of ISSTA*, 2008.
- 55. B. Wegbreit. The synthesis of loop predicates. Communications of the ACM, 17(2):102-112, 1974.
- 56. T. Xie, D. Marinov, W. Schulte, and D. Notkin. Symstra: A framework for generating object-oriented unit tests using symbolic execution. In Proc. of the 11th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS), 2005.
- 57. T. Yavuz-Kahveci and T. Bultan. Automated verification of concurrent linked lists with counters. In G. P. M. Hermenegildo, editor, Proceedings of the 9th International Static Analysis Symposium, volume 2477 of Lecture Notes in Computer Science. Springer-Verlag, 2002.
- 58. G. Yorsh, T. Ball, and M. Sagiv. Testing, abstraction, theorem proving: better together! In *Proc. ISSTA*, 2006.